

LISTING OF THE CLAIMS:

Claims 1-20 (Cancelled).

21. (New) A method for building globally optimal configurations of one or more clusters of networked resources, for bringing them on-line in a systematic manner, given a set of resources, resource groups, their current states, dependencies, preferences, constraints, events, and policies, including concurrently discovered resources and resource groups, availability and unavailability of the resources and resource groups for being able to go on-line, determining their dynamic dependencies and configuration information such as load serving capacity and quality of service, cluster policies and changes thereof, as applied to the network of resources and resource groups, at cluster initialization and dynamically during cluster operation, for supporting a seamless startup and shutdown of resources and resource clusters according to the current policies, multiple demands on the cluster resources, current status of the cluster resources, system events, and their effects on the priority of the services provided by the cluster, separating the said dependencies, constraints, events, and policies into a static or an occasionally changing rules and objectives group and a dynamically changing cluster events and policies group, and separating the networked resources, resource groups, cluster configurations into static and dynamically changing groups and taking a snapshot of the said groups only when needed to build the said optimal cluster configurations.

22. (New) A method according to Claim 21, comprising the further steps of:

continuously monitoring cluster-wide events and comparing the current cluster state with a desired state, and whenever there is a discrepancy between said current and desired states, realigning the cluster resources, including the step of issuing commands to the cluster resources to bring about the realigning;

providing a group of cluster services, including:

- i) a persistent cluster registry to store and retrieve the configuration of the cluster resources,
- ii) topology services for detecting node and communication adapter failures,
- iii) reliable messaging for selected communications between a central resource and all other resources, and
- iv) a group services facility for electing one of the resources as the central resource at cluster initialization and whenever an existing central resource is unable to provide the services thereof,

delivering events to a coordinator, said coordinator combining said events with said rules and objectives to arrive at a response to said events;

translating the response into commands to the resources, each of the commands containing all the state needed for execution of the command by a manager of one of the resources, including the step of issuing the commands in a partial order given by said dependencies; and

not sending out a new command until the leader resource is aware of a positive outcome of the commands that the execution of said new command depends on.

23. (New) A method according to Claim 22, wherein:

said coordinator ensures that globally-optimal solutions get deployed in the cluster in response to events in the cluster; and

all events and command feedback are directed to said coordinator.

24. (New) A method according to Claim 21, comprising the further steps of:

providing an optimizer module for computing a globally optimal solution based on said constraints and to current state of the cluster;

using the optimizer for recomputing the globally optimal solution whenever an objective value of a deployed solution is below a certain value as compared to a proposed solution, including the step of feeding back to the optimizer an artificially generated event that forces the optimizer to recompute the global solution;

providing the optimizer with a snapshot of the current state of the cluster;

wherein the step of using the optimizer for recomputing the globally optimal solution includes the step of said optimizer, given said snapshot, proposing an approximately optimal cluster configuration that takes into account said current state of the cluster and long-term objectives defined for the cluster.

25. (New) Apparatus for building globally optimal configurations of one or more clusters of networked resources, for bringing them on-line in a systematic manner, given a set of resources, resource groups, their current states, dependencies, preferences, constraints, events, and policies, including concurrently discovered resources and resource groups, availability and unavailability of the resources and resource groups for being able to go on-line, determining their dynamic dependencies and configuration information such as load serving capacity and quality of service, cluster policies and changes thereof, as applied to the network of resources and resource groups, at cluster initialization and dynamically during cluster operation, for supporting a seamless startup and shutdown of resources and resource clusters according to the current policies, multiple demands on the cluster resources, current status of the cluster resources, system events, and their effects on the priority of the services provided by the cluster, separating the said dependencies, constraints, events, and policies into a static or an occasionally changing rules and objectives group and a dynamically changing cluster events and policies group, and separating the networked resources, resource groups, cluster configurations into static and dynamically changing groups and taking a snapshot of the said groups only when needed to build the said optimal cluster configurations.

26. (New) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for building globally optimal configurations of one or more clusters of networked resources, for bringing them on-line in a systematic manner, given a set of resources, resource groups, their current states, dependencies, preferences, constraints, events, and policies, including concurrently discovered resources and resource groups, availability and unavailability of the resources and resource groups for being able to go on-line, said method steps comprising determining their dynamic dependencies and configuration information such as load serving capacity and quality of service, cluster policies and changes thereof, as applied to the network of resources and resource groups, at cluster initialization and dynamically during cluster operation, for supporting a seamless startup and shutdown of resources and resource clusters according to the current policies, multiple demands on the cluster resources, current status of the cluster resources, system events, and their effects on the priority of the services provided by the cluster, separating the said dependencies, constraints, events, and policies into a static or an occasionally changing rules and objectives group and a dynamically changing cluster events and policies group, and separating the networked resources, resource groups, cluster configurations into static and dynamically changing groups and taking a snapshot of the said groups only when needed to build the said optimal cluster configurations.